

# Real-Time Predictive Modeling and Robust Avoidance of Pedestrians with Uncertain, Changing Intentions

Sarah Ferguson, Brandon Luders, Robert C. Grande, and Jonathan P. How

Aerospace Controls Laboratory  
Massachusetts Institute of Technology  
Cambridge MA 02139

**Abstract.** To plan safe trajectories in urban environments, autonomous vehicles must be able to quickly assess the future intentions of dynamic agents. Pedestrians are particularly challenging to model, as their motion patterns are often uncertain and/or unknown *a priori*. This paper presents a novel changepoint detection and clustering algorithm that, when coupled with offline unsupervised learning of a Gaussian process mixture model (DPGP), enables quick detection of changes in intent and online learning of motion patterns not seen in prior training data. The resulting long-term movement predictions demonstrate improved accuracy relative to offline learning alone, in terms of both intent and trajectory prediction. By embedding these predictions within a chance-constrained motion planner, trajectories which are probabilistically safe to pedestrian motions can be identified in real-time. Hardware experiments demonstrate that this approach can accurately predict pedestrian motion patterns from onboard sensor/perception data and facilitate robust navigation within a dynamic environment.

**Keywords:** pedestrian modeling, intent prediction, Gaussian processes, probabilistic path planning, autonomous vehicles

## 1 Introduction

Autonomous vehicles operating in urban environments must be able to quickly assess the future behavior of nearby agents in order to plan safe trajectories. A major challenge in navigating such environments is the limited ability to accurately anticipate the intents of dynamic agents, as their internal state and desires are not directly observable. Dynamic agents exhibit uncertain in both their intent and the trajectory motion pattern associated with each intent. Pedestrians present particular technical challenges in the generation of long-term predictions, specifically (i) the demonstration of many unique behaviors, which may not have been previously observed; and (ii) instantaneous changes in motion behavior following changes in intent.

This paper addresses these challenges by proposing a novel framework for long-term predictions of pedestrian motion, capable of learning new motion patterns and quickly responding to changes in intent as they are observed online.

Such a predictive model is necessary because naively modeling the uncertainty in dynamic environments can lead to rapidly increasing uncertainty, which is typically prohibitive for safe navigation [22]. The preferred approach in the literature, also used here, assumes that factors influencing pedestrian motion (such as internal state and intent) are reflected by their trajectories. Therefore, various data-driven approaches learn typical motion patterns from observed training trajectories to enable predictions of future state.

The most common approaches are based on the Markov property. This set of approaches includes hidden Markov models, in which the hidden state is pedestrian intent [4,14,24]; growing hidden Markov models to allow for online learning [23]; and partially observable Markov decision processes to choose actions based on a distribution over pedestrian intents [2]. Because the future state prediction is dependent only on the current state, these approaches are quick to react to changes in intent. However, for relatively infrequent changes in intent, the Markov assumption can be overly restrictive, as it prevents these algorithms from becoming more certain of pedestrian intent with additional observations.

Gaussian process (GP) approaches have been demonstrated to be well-suited for modeling pedestrian motion patterns, as they perform well with noisy observations and have closed-form predictive uncertainty [6,21,19]. Additionally, recent work using GP mixture models enables predictions that account for both intent and trajectory uncertainty [1]. Both sets of approaches use the entire observed trajectory in the prediction of future state, such that certainty in demonstrated intent tends to converge over time. Therefore, when changes in intent occur, these approaches are much slower to detect a change than Markov-based approaches. Additionally, existing GP classification approaches are too slow for online learning of previously unobserved behavior patterns.

This paper proposes a novel changepoint detection and clustering algorithm which retains the trajectory prediction accuracy of existing GP approaches while expanding their capabilities. Coupled with offline unsupervised learning of a Gaussian process mixture model (DPGP) [13], this approach enables quick detection of changes in intent and online learning of motion patterns not seen in prior training data. The resulting long-term movement predictions demonstrate improved accuracy relative to offline learning alone in both intent and trajectory prediction. These predictions can also be used within a chance-constrained motion planner [15] to identify probabilistically safe trajectories in real-time. In experimental results, the proposed algorithm is used to predict the motion of multiple dynamic agents detected from a variety of onboard and external sensors, enabling an autonomous rover to robustly navigate dynamic environments.

## 2 Preliminaries

### 2.1 Motion Patterns and Modeling

A pedestrian trajectory is represented as a set of observed  $xy$ -locations  $(x_1^i, y_1^i), (x_n^i, y_n^i), \dots, (x_{L^i}^i, y_{L^i}^i)$ , where  $L^i$  is the total length of the trajectory  $t^i$  of pedestrian  $i$ . Trajectories are not necessarily of the same length, and the time steps

between each observation may be irregular. A motion pattern is defined as a mapping from  $xy$ -locations  $(x^i, y^i)$  to a distribution over trajectory derivatives  $\left(\frac{\Delta x^i}{\Delta t}, \frac{\Delta y^i}{\Delta t}\right)$ , resulting in a velocity flow-field in  $x - y$  space. This approach is therefore independent of the lengths and discretization of the trajectories.

This work uses Gaussian processes (GP) for its motion pattern models. Although GPs have a significant mathematical and computational cost, they generalize well to regions of sparse data while avoiding the problem of over fitting in regions of dense data. GP models are extremely robust to unaligned, noisy measurements and are well-suited for modeling the continuous paths underlying potentially non-uniform time-series samples of pedestrian locations [19].

The GP serves as a non-parametric form of interpolation between the discrete trajectory measurements. Specifically, given an observed  $(x, y)$  location, the GP predicts the trajectory derivatives at that location. The standard squared exponential covariance function describes the correlation between trajectory derivatives at two points  $(x, y)$  and  $(x', y')$ . The mean trajectory derivative functions  $E[\frac{\Delta x^i}{\Delta t}, \frac{\Delta y^i}{\Delta t}] = \mu_x(x, y)$  and  $E[\frac{\Delta y^i}{\Delta t}, \frac{\Delta x^i}{\Delta t}] = \mu_y(x, y)$  are implicitly initialized to zero for all  $xy$  locations.

The resulting motion model is defined as a finite mixture of the  $M$  learned motion patterns weighted by their prior probabilities. The finite mixture model probability of the  $i$ th observed trajectory  $t^i$  is

$$p(t^i) = \sum_{j=1}^M p(b_j)p(t^i|b_j), \quad (1)$$

where  $b_j$  is the  $j$ th motion pattern and  $p(b_j)$  is its prior probability. The prior  $p(b_j)$  is initialized to be proportional to the number of trajectories in motion model  $j$ , but is updated after each prediction with the previous posterior probability which incorporates the GP likelihood. The number of motion patterns  $M$  is learned offline via an automated clustering process and incremented as new behavior patterns are identified online. This motion model has been previously presented in [1,13].

Future pedestrian trajectories are predicted for each motion pattern using the approach of [7,5]. This approach provides a fast, analytic GP approximation specifying possible future pedestrian locations  $K$  time steps ahead, while incorporating uncertainty in previous predictions at each time step.

## 2.2 Batch Learning of Motion Patterns

It is expected that observed pedestrian trajectories will demonstrate a variety of qualitatively different behaviors. These behavior motion patterns are learned from an input set of unlabeled trajectories by DPGP, a Bayesian nonparametric clustering algorithm that automatically determines the most likely number of clusters without *a priori* information [13]. This section reviews the DPGP algorithm, which is used in this work to cluster observed pedestrian trajectories into representative motion patterns in batch.

The DPGP algorithm models motion patterns as Gaussian processes, as described in Sect. 2.1. A Dirichlet process (DP) mixture model potentially allows for an infinite mixture of motion patterns, where the DP concentration parameter  $\alpha$  controls the probability of new cluster formation. A smaller  $\alpha$  enforces the expectation that there are a few motion patterns that pedestrians tend to exhibit; therefore, trajectories are more likely to fit existing clusters than to form new ones.

Because exact inference over the space of GPs and DPs is intractable, samples are drawn from the posterior over motion models to train the model. At each iteration, the DP hyperparameter  $\alpha$  is resampled using Gibbs sampling techniques and the GP hyperparameters for the  $j$  behavior patterns  $\theta_{x,j}^{GP}, \theta_{y,j}^{GP}$  are set to their maximum likelihood values using gradient ascent. The probability that trajectory  $t^i$  will be assigned to an existing motion pattern is

$$p(z_i = j | t^i, \alpha, \theta_{x,j}^{GP}, \theta_{y,j}^{GP}) \propto l(b_j; t^i) \left( \frac{n_j}{N - 1 + \alpha} \right), \quad (2)$$

and the probability that trajectory  $t^i$  will be assigned to a new motion pattern is

$$p(z_i = M + 1 | t^i, \alpha) \propto \int l(b_j; t^i) d\theta_{x,j}^{GP} d\theta_{y,j}^{GP} \left( \frac{\alpha}{N - 1 + \alpha} \right), \quad (3)$$

where  $l(b_j; t^i)$  is the likelihood that trajectory  $t^i$  is assigned to motion pattern  $b_j$ ,  $n_j$  is the number of trajectories currently assigned to  $b_j$ , and  $N$  is the total number of trajectories. The likelihoods  $l(b_j; t^i)$  can be determined from the GP trajectory derivatives.

### 2.3 Motion Planning

Motion planning for autonomous vehicles is executed via chance-constrained rapidly-exploring random trees (CC-RRT), which can efficiently identify trajectories with guaranteed minimum bounds on constraint satisfaction probability under internal and/or external uncertainty [15]. The primary objective is to plan and execute a motion plan directing the vehicle to reach some goal region, while ensuring the non-convex state constraints  $x_t \in \mathcal{X}_t$  are probabilistically satisfied. This is represented via path-wise and time-step-wise chance constraints

$$\mathbb{P} \left( \bigwedge_t x_t \in \mathcal{X}_t \right) \geq \delta_p, \quad \mathbb{P}(x_t \in \mathcal{X}_t) \geq \delta_s, \quad \forall t, \quad (4)$$

respectively, where  $\mathbb{P}(\cdot)$  denotes probability,  $\bigwedge$  represents a conjunction over the indexed constraints, and  $\delta_s, \delta_p \in [0.5, 1]$  are chosen by the user. The feasible state space  $\mathcal{X}_t$  consists of a convex environment containing multiple convex, polytopic obstacles to be avoided. It is assumed that the shape and orientation of these obstacles is known, but their placement may be uncertain and/or dynamic.



The CC-RRT algorithm samples a tree of dynamically and probabilistically feasible trajectories through the environment, rooted at the vehicle’s current state. All trajectories added to the tree must satisfy (4), which CC-RRT evaluates by leveraging the trajectory-wise constraint checking of sampling-based algorithms to efficiently compute risk bounds [15].

In this work, detected pedestrians are modeled as dynamic obstacles, with both intent and trajectory uncertainty as represented by (1). This model provides a likelihood and time-parameterized uncertainty distribution for each behavior of each pedestrian obstacle. The CC-RRT formulation has been recently expanded to guarantee probabilistically robust avoidance of dynamic obstacles with uncertain intentions, making it suitable for robust avoidance of pedestrian models. In particular, if each dynamic pedestrian behavior is treated as a separate obstacle whose risk bound is scaled by its likelihood, then all probabilistic feasibility guarantees are maintained [1].

### 3 Changepoint Detection

To effectively anticipate the motion of pedestrians, this paper proposes a framework which can perform online classification of observed trajectories, in addition to learning common pedestrian trajectories from batch data. Because agile dynamic agents such as pedestrians may exhibit new behaviors or mid-trajectory changes in intent, this problem is framed in the context of changepoint detection.

This work utilizes a variation of the generalized likelihood test (GLR) [3] to perform changepoint detection. The basic GLR algorithm detects changes by comparing a windowed subset of data to a null hypothesis. If the maximum likelihood statistics of the windowed subset differ from the null hypothesis significantly, the algorithm returns that a changepoint has occurred.

The proposed changepoint detection algorithm is given in Algorithm 1. At each time step, given Gaussian process  $GP_w$ , the algorithm creates a new GP ( $GP_S$ ) with the same hyperparameters, but using a windowed data subset  $S$  of size  $m_S$  (lines 2–4). Although  $m_S$  is domain specific, the algorithm is fairly robust to its selection;  $m_S \approx 10 - 20$  works well for most applications.

The algorithm then calculates the joint likelihood of the set having been generated from the current GP model (the null hypothesis  $H_0$ ) and the new  $GP_S$  ( $H_1$ ). At each step, the normalized log-likelihood ratio test (LRT) is computed as

$$L(y) = \frac{1}{m_s} (\log P(S | H_1) - \log P(S | H_0)). \quad (5)$$

For a GP, the log likelihood of a subset of points can be evaluated in closed form as

$$\log P(y | x, \Theta) = -\frac{1}{2} (y - \mu(x))^T \Sigma_{xx}^{-1} (y - \mu(x)) - \log |\Sigma_{xx}|^{1/2} + C,$$

where  $\mu(x)$  is the mean prediction of the GP and

$$\Sigma_{xx} = K(x, x) + \omega_n^2 I - K(X, x)^T (K(X, X) + \omega_n^2 I)^{-1} K(X, x)$$

---

**Algorithm 1** Likelihood Ratio Test

---

- 1: **Input:** Set of points  $\mathcal{S}$ , Working model  $GP_w$
  - 2:  $l_1 = \log p(\mathcal{S} \mid GP_w)$
  - 3: Create new GP  $GP_S$  from  $\mathcal{S}$
  - 4:  $l_2 = \log p(\mathcal{S} \mid GP_S)$
  - 5: Calculate LRT  $L_i(y) = \frac{1}{m_S} (l_2 - l_1)$
  - 6: Calculate average of last  $m$  LRT:  
$$L_m = \frac{1}{m} \sum_{j=i-m}^i L_j(y)$$
  - 7: Calculate average of LRT after changepoint:  
$$L_{ss} = \frac{1}{i-m-1} \sum_{j=1}^{i-m-1} L_j(y)$$
  - 8:  $i = i + 1$
  - 9: **return**  $L_m - L_{ss} \geq \eta$
- 

is the predictive variance of the GP plus the measurement noise. The first term of the log-likelihood accounts for the deviation of points from the mean, while the second accounts for the relative certainty (variance) in the prediction.

Algorithm 1 uses the LRT to determine if the maximum likelihood statistics (mean and variance) of  $GP_S$  differ significantly from the null hypothesis, indicating that the points in  $\mathcal{S}$  are more unlikely to have been generated from the model  $GP_w$ . In particular, the average over the last  $m$  LRT values (line 6) is compared to the nominal LRT values seen up until this point (line 7). If the difference of these two values exceeds some value  $\eta$ , the algorithm returns false, indicating that this generating model does not fit the data. The value  $\eta$  can be determined based on the probability of false alarms and maximum allowed error [8].

The LRT algorithm is quite robust in practice, based on the following intuition. If the points in  $\mathcal{S}$  are anomalous simply because of output noise, then the new GP model created from these points will on average be similar to the current model. Additionally, the joint likelihood given the new model will not be substantially different from that of the current model. However, if the points are anomalous because they are drawn from a new process, then the resulting GP model will on average be substantially different from the current model, yielding a higher joint likelihood of these points. Lastly, instead of making a decision on a single LRT, the last  $m$  LRT's are averaged and compared to the average LRT values seen since the last changepoint. In practice, the LRT may have some offset value due to modeling error. Looking at the difference between the last  $m$  values and the average LRT values makes the algorithm robust to this problem.

## 4 Changepoint-DPGP

In order to generate accurate and computationally efficient predictions of pedestrian motion, the proposed Changepoint-DPGP algorithm seeks to identify new behaviors online and detect changes in intent given typical pedestrian behaviors learned from batch data. The key idea behind this algorithm is to perform online

---

**Algorithm 2** Changepoint-DPGP

---

```
1: Input: Set of previous behavior models  $\mathcal{GP} = \{GP_1, \dots, GP_N\}$ 
2: while Input/Output  $\langle x_t, y_t \rangle$  available do
3:   Add  $\langle x_t, y_t \rangle$  to  $\mathcal{S}$ 
4:   Call Algorithm 3
5:   if  $\mathcal{M}_{t-1} \cap \mathcal{M}_t = \emptyset$  then                                     # Change in intent detected
6:     Reinitialize priors
7:   end if
8:   if  $\mathcal{M}_t = \emptyset$  then                                             # New behavior detected
9:     Initialize new model  $GP_n$ 
10:  else
11:    Predict according to Sect. 2.1
12:  end if
13:  if  $\mathcal{M}_t \neq \emptyset$  then
14:     $\mathcal{M}_t = \mathcal{M}_{t-1} \cap \mathcal{M}_t$ 
15:  end if
16: end while
17: if  $GP_n$  is initialized then
18:   Add  $\langle x_{0:T}, y_{0:T} \rangle$  to  $GP_n$ 
19:   Add  $GP_n$  to set of current models  $\mathcal{GP}$ 
20: end if
```

---

---

**Algorithm 3** Compare to Current Models

---

```
1: Input: Set of current behavior models  $\mathcal{GP} = \{GP_1, \dots, GP_N\}$ 
2: Initialize representative model set  $\mathcal{M}_t$ 
3: for Each  $GP_j \in \mathcal{GP}$  do
4:   Call Algorithm 1 with inputs  $\mathcal{S}, GP_j$ 
5:   if Algorithm 1 returns true then
6:     Add  $GP_j$  to  $\mathcal{M}_t$ 
7:   end if
8: end for
```

---

classification of a sliding window of trajectory segments, and detect changepoints or new behavior models according to changes in the current classification.

The Changepoint-DPGP algorithm is detailed in Algorithm 2. The algorithm begins with an initial set of learned behavior motion models  $\mathcal{GP}$ , obtained from running the DPGP algorithm on batch training data. As new data points are received, they are added to a sliding window  $\mathcal{S}$  of length  $m_s$ . After creating a new model  $GP_S$  from the points in  $\mathcal{S}$ , the LRT is computed for  $GP_S$  and for each model  $GP_j$  in the current model set  $\mathcal{GP}$ . This process determines if the points in  $\mathcal{S}$  are statistically similar to those in the model  $GP_j$ , subject to the predetermined threshold  $\eta$ .

In order to detect changepoints, the algorithm maintains the set of models  $\mathcal{M}_t$  that the points in  $\mathcal{S}$  fit into at each time step, representative of the current classification of those points. Because the behavior patterns may overlap (e.g. the blue/green and red/teal behavior patterns in Fig. 1(b)), a single classification

cannot be guaranteed, necessitating the maintenance of a model set. Change-points occur when the classification changes, i.e. when the current classification  $\mathcal{M}_t$  and previous classification  $\mathcal{M}_{t-1}$  share no common models. Additionally, the current classification is reset at each timestep to be the intersection of the current and previous classification sets, assuming that the current classification is not empty.

To illustrate this method, consider a pedestrian crossing a crosswalk by following the green behavior pattern in Fig. 1(b). Until the pedestrian reaches the crosswalk,  $\mathcal{M}_t = \{B, G\}$ . Once the pedestrian enters the crosswalk, their classification becomes  $\mathcal{M}_t = \{G\}$ . A changepoint should not be detected at this stage, as the pedestrian is committing to the green behavior rather than changing their intent. However, if the pedestrian switched to the teal behavior after entering the crosswalk, this would represent a change in intent. The classification for three successive timesteps would become  $\mathcal{M}_{t-2} = \{G\}$ ,  $\mathcal{M}_{t-1} = \{G, T\}$ ,  $\mathcal{M}_t = \{T\}$  and no changepoint would be detected if  $\mathcal{T}_t$  was not reset.

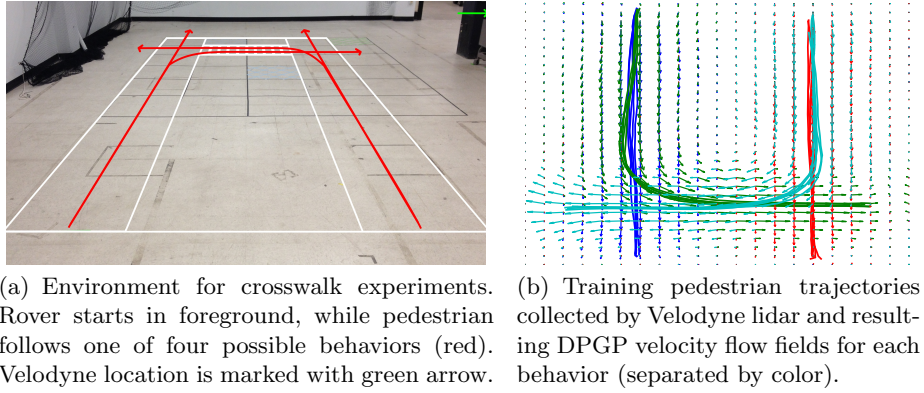
The predictive component of this algorithm is decoupled from classification. In general, the future state distribution is computed as described in Sect. 2. However, if at any point  $\mathcal{M}_t$  is empty, this indicates that the current model set  $\mathcal{GP}$  is not representative for the points in  $\mathcal{S}$ , so a new behavior must be created. The algorithm waits until the entire new trajectory has been observed to create the new behavior pattern, generating predictions according to a simple velocity propagation model until the model set becomes representative. In practice, any reasonable predictive model can be used at this stage, as no information on the anomalous agent’s current behavior.

If the training data contains trajectories with changes in intent, DPGP will learn unique behavior patterns for each trajectory containing a changepoint, as the entire trajectory is considered for classification. To obtain a representative set of behavior patterns, the Changepoint-DPGP algorithm can be used offline to reclassify these trajectories by segmenting them at the changepoint. To do so, Algorithm 2 is first called with  $\mathcal{GP}$  containing those behavior patterns with more than  $k_{min}$  trajectories and data  $\langle x_t, y_t \rangle$  from trajectories in the remaining behavior patterns not in  $\mathcal{GP}$ . At line 6 and at the end of Algorithm 2, the trajectory segment seen since the last changepoint is classified into the most likely behavior pattern. The intuition behind this modification is that changes in intent are agent-specific; therefore, behavior patterns containing these trajectories are not representative of global behaviors caused by the environment.

As the Changepoint-DPGP algorithm is called for each dynamic agent, with learned motion models in  $\mathcal{GP}$  possibly shared among all agents, Algorithms 2-3 can be parallelized to speed up computation if desired.

## 5 Results

This section presents empirical results which evaluate Changepoint-DPGP on real-world problem domains of varying complexity. The prediction results demonstrate that prior observations of pedestrian motion can be used to learn accurate



**Fig. 1.** Environment setup and pedestrian data for crosswalk experiments.

behavior models. These models are applied to real-time observations to make accurate, long-term predictions of complex motion behavior, beyond what could be predicted from the observations themselves (e.g., bearing and speed). The planner is then demonstrated to select safe paths which are risk-aware with respect to possible pedestrian intentions, their likelihood, and their risk of interaction with the host vehicle. The final version of this paper will include more comprehensive results and analysis.

### 5.1 Pedestrian Crosswalk

Consider the scenario in Fig. 1, in which an autonomous rover travels along a street flanked by two sidewalks and must safely pass through a pedestrian crosswalk. Pedestrians have four possible behaviors (red) corresponding to which sidewalk they are traversing, and whether they choose to use the crosswalk.

A Pioneer 3-AT rover is used as the autonomous vehicle in this and subsequent experiments. Its payload includes a SICK LMS-291 lidar for onboard pedestrian detection and an Intel Core i5 laptop with 6GB RAM for computation. The online perception, planning, and control algorithms described in this paper are executed on this laptop via the Robotic Operating System (ROS) [18]. Dynamic obstacle detections and autonomous vehicle state are fed to a real-time, multi-threaded Java application, which executes the chance-constrained planner (Sect. 2.3) to generate safe paths. A pure pursuit controller [17] generates acceleration commands to follow this set of waypoints. High-fidelity localization is provided for the rover via motion-capture cameras [10].

Three trajectory prediction algorithms are evaluated in this experiment: changepoint-DPGP, DPGP, and a goal-directed approach using hidden Markov models (HMM) [16]. The hidden states of the HMM are pedestrian goals, learned via Bayesian nonparametric inverse reinforcement learning with an approximation to the action likelihood specifying that pedestrians head directly towards

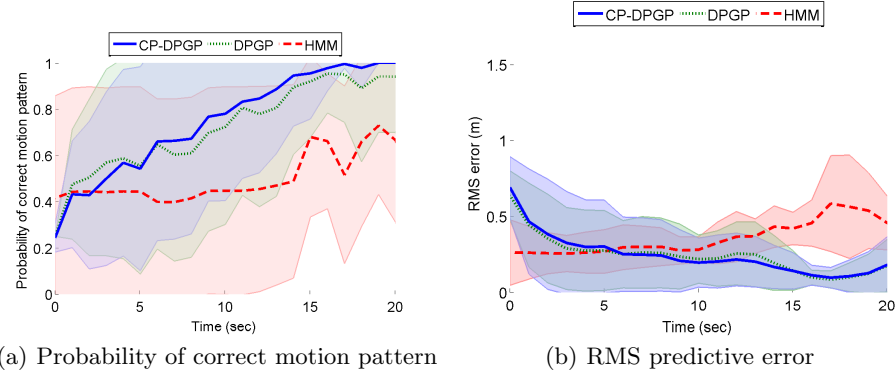
goal locations [16]. This motion model assumes that each pedestrian heads directly toward their intended goal at some preferred speed with an uncertainty distribution over heading and velocity, as used by [9,2,11] among others.

Unless otherwise noted, all three algorithms were trained on five trajectories from each of the four behavior patterns in Fig. 1(a). Each trajectory was collected by taking observations of an actual pedestrian traversing the environment, as observed by a Velodyne HDL-32E lidar at the location marked in green in Fig. 1(a). Pedestrians are identified from the raw Velodyne returns both offline and online using Euclidean clustering [20]. Fig. 1(b) shows the training trajectories used in this experiment.

Fig. 2 considers the baseline case in which (i) neither the training nor test data exhibits any mid-trajectory changes in intent; and (ii) the test data does not exhibit any behaviors unseen in the training data. In these results, each algorithm is tested on five trajectories from each of the four behavior patterns (Fig. 1(a)). Fig. 2(a) displays the probability each algorithm has assigned to the correct motion pattern given the observation trajectory, averaged across all 20 trials as a function of time elapsed. This metric measures the ability of each algorithm to identify the correct pedestrian intentions. The likelihoods of each motion pattern serve as the intent prediction for the GP-based approaches, with the prior probability (time = 0) based on the fraction of training trajectories for each motion pattern. The most likely state distribution, calculated via the forward algorithm and Markov chain propagation, describes the predicted intent for the HMM approach. Fig. 2(b) displays the root mean square (RMS) error between the true pedestrian position and the mean predicted position (Sect. 2.1), averaged across all 20 trials as a function of time elapsed. This metric captures overall prediction accuracy subject to both intent and path uncertainty.

The Markov property prevents the HMM approach from converging to the correct motion pattern, as the observations of current state alone are not sufficient in the case of noisy observations (Fig. 2(a)). As a result, its RMS error tends to increase over time. On the other hand, both GP approaches exhibit convergence in the probability of the correct motion pattern as new observations are made, which improves RMS predictive error as well. The performance of changepoint-DPGP and DPGP is very similar, as is expected in the absence of changepoints and new behaviors.

Next, each algorithm is tested on five trajectories which demonstrate a change in pedestrian intentions, or a changepoint. In these trajectories, the pedestrian begins to traverse the crosswalk, but reverses direction after 18 seconds. Fig. 3 shows the evolution of the correct likelihood and RMS error for each algorithm in this scenario, averaged across the trajectories. Both DPGP and Changepoint-DPGP converge on the correct behavior prior to the changepoint (Fig. 3(a)), while HMM performance is relatively unchanged compared to Fig. 2(b). As the change in pedestrian intention takes place, both GP-based algorithms initially drop to zero probability, as expected. However, DPGP accuracy remains poor beyond the changepoint, leading to the largest RMS errors of all algorithms (Fig. 3(b)). Because DPGP relies on the entire observation history, its predictions

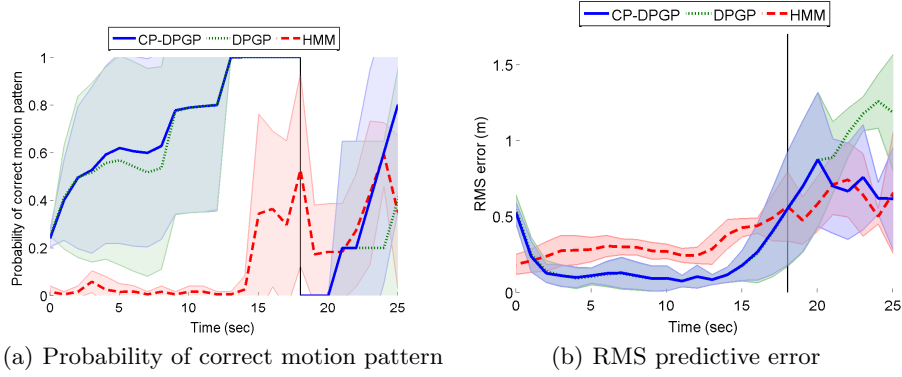


**Fig. 2.** Prediction accuracy of each algorithm for the baseline case of the pedestrian crosswalk scenario. All results are averaged over 20 trials as a function of time elapsed, with error bars representing standard deviation.


are slow to recognize the change, leading to worse performance. On the other hand, Changepoint-DPGP is able to selectively update the observation history considered in the likelihood computation given changes in intent, enabling it to achieve better accuracy than DPGP after the changepoint (Fig. 3(a)). As a result, Changepoint-DPGP yields the lowest average overall RMS error of all algorithms tested (Fig. 3(b)).

Changepoint-DPGP also demonstrates the best relative prediction accuracy when considering anomalous/new behavior patterns. In this scenario, algorithms are trained on only three of the four possible behaviors (red, blue, green in Fig. 1(b)), then tested on five trajectories from the fourth behavior (teal in Fig. 1(b)). The teal behavior deviates from the previously-observed red behavior approximately 9 seconds into the trajectory. Fig. 4 shows the evolution of the RMS error for each algorithm in this scenario. At 9 seconds into the experiment, when the pedestrian behavior begins to deviate from anything observed in training data, the prediction of both HMM and DPGP begins to steadily increase. On the other hand, Changepoint-DPGP successfully identifies the new behavior and reclassifies subsequent trajectories. Thus it exhibits behavior similar to the baseline case, in which predictive error decreases as the probability of the correct motion pattern converges. Overall, Changepoint-DPGP predictive error is reduced by 62% compared to DPGP. This demonstrates the strength of Change-DPGP in cases where the training data is not representative of the observed motion patterns, e.g. due to short periods of data collection.


Finally, experiments have demonstrated that predictive results from the proposed Changepoint-DPGP enable an autonomous rover (Fig. 5) to safely avoid collision in closed-loop. Fig. 6 gives snapshots of a representative interaction between a pedestrian and the autonomous rover. In this experiment, the rover localizes itself using motion-capture data, and identifies and localizes pedestri-



**Fig. 3.** Prediction accuracy of each algorithm for pedestrian crosswalk scenario, subject to pedestrian change in intentions at time = 18 s.



**Fig. 4.** RMS error for pedestrian crosswalk scenario, subject to trajectories not observed in training data.



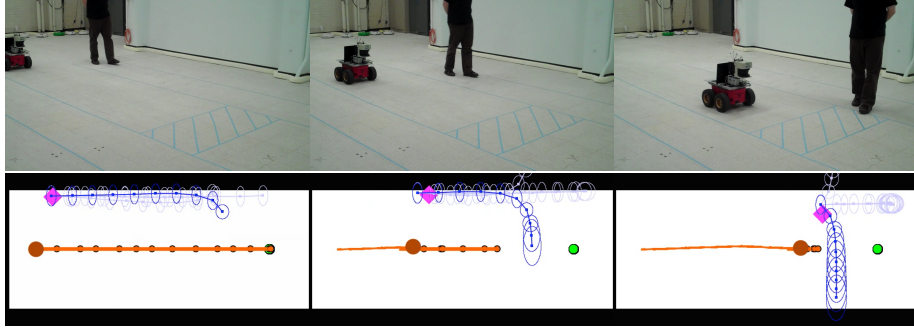
**Fig. 5.** Rover used in closed-loop motion planning experiments.

ans exclusively from its onboard SICK lidar. The rover then uses the motion pattern predictions within a CC-RRT probabilistic motion planner (Sect. 2.3) to generate safe trajectories. Initially, the planner generates a path directly to the goal, as the pedestrian is projected to remain on the sidewalk (Fig. 6, left). Once the predictions indicate that the pedestrian is likely to cross, the planner adjusts its plan to terminate prior to the crosswalk (Fig. 6, center). Once the pedestrian begins to cross (Fig. 6, right), the rover comes to a stop, waiting for the crosswalk to clear before safely proceeding to the goal.

## 5.2 Multiple Dynamic Vehicles

In this experiments, the autonomous rover must travel through a continuous sequence of waypoints while safely avoiding one or more dynamic robot obsta-





**Fig. 6.** Snapshots of a representative rover/pedestrian interaction, in which the autonomous rover safely avoids a pedestrian at the crosswalk. The autonomous rover (brown) seeks to reach the goal location (green) while avoiding the pedestrian (magenta). The rover’s current plan/history (orange) and predicted pedestrian behaviors (blue; darker shades indicate higher likelihoods) are shown.

cles (referred to below simply as “robots”), using Changepoint-DPGP to predict robot motion and CC-RRT to generate safe trajectories. Each waypoint is represented as a goal region (Fig. 7) within the RAVEN motion-capture environment [10]; the rover is provided with the next waypoint once the previous goal region is reached. The iRobot Create vehicle [12] is used as the dynamic robot obstacle in these experiments. It continuously executes one of four cyclical, counter-clockwise motion patterns within the testbed (Fig. 7, first snapshot). State information for both the autonomous rover and iRobot is provided by the RAVEN motion-capture cameras.

Fig. 7 illustrates how Changepoint-DPGP behaves when a new behavior pattern is seen repeatedly. The rover has received training trajectories for behavior 1 offline, but a single dynamic robot actually executes behavior 3 online. Initially, the only known behavior is behavior 1, so the autonomous rover is certain that the robot will pass between it and the goal and modifies its path accordingly. At 8 seconds, changepoint detection recognizes that the robot is executing a new behavior. Predictions are then generated assuming that the robot will continue at its current velocity with increased, linearly-scaling uncertainty. The planner modifies its planned paths to the goal to reflect this shift in perceived behavior.

After 85 seconds, the algorithm has learned the entire trajectory that it has just observed, as a new behavior. As the robot begins its second cycle, it still assigns the highest likelihood to behavior 1, based on the prior distribution of observed training and test trajectories still favoring behavior 1. However, behavior 3 is now included as an additional behavior prediction. By 97 seconds, the algorithm is fully confident that the robot is executing behavior 3, and shifts its likelihoods accordingly. The predictions for the new behavior accurately reflect the trajectory executed by the robot with reduced uncertainty. Based on this reduced uncertainty, the planner knows that the robot will turn before

intersecting with the autonomous rover’s planned path, and thus continues to execute that path. This scenario was executed for 2.5 minutes with no collisions.

Changepoint-DPGP has also been demonstrated to operate safely in real time with multiple dynamic obstacles; the final version of this paper will provide greater detail on this multi-robot examples. Fig. 8 provides images from a representative baseline scenario, in which two dynamic robots are trained on and exhibit all four possible behaviors online. The scenario was executed for 2 minutes with no collisions.

## 6 Conclusions

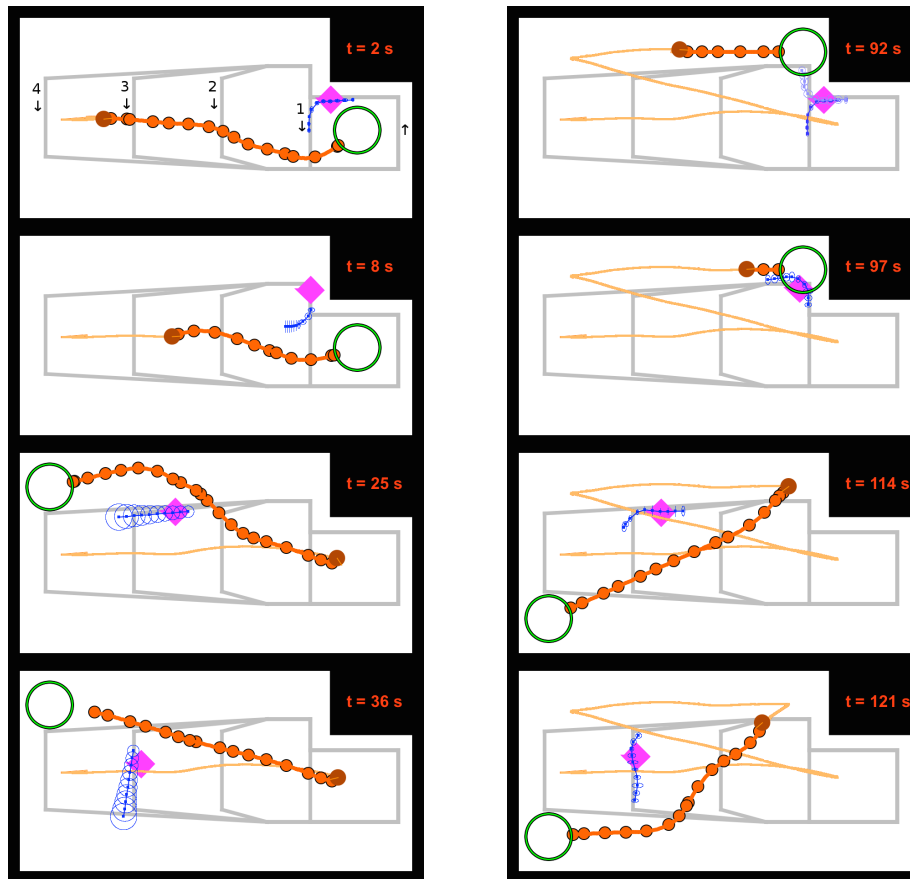
This paper has developed a framework for long-term trajectory prediction and robust collision avoidance of pedestrians and other dynamic agents in real-time, even when these agents exhibit previously unobserved behaviors or changes in intent. A key contribution is the Changepoint-DPGP algorithm, which uses a non-Bayesian likelihood ratio test for efficient online classification and detection of changepoints. This algorithm is able to learn new behavior patterns online and quickly detect and react to changepoints, capabilities currently not present for other predictive algorithms that use GP motion models. As demonstrated in real-time simulation results, these capabilities significantly improve prediction accuracy relative to existing methods. Preliminary hardware results show that the framework can accurately predict motion patterns of dynamic agents from various sensor data and perform robust navigation.

## Acknowledgments



Research funded by Ford Motor Company (James McBride, Ford Project Manager).

## References

1. Aoude, G.S., Luders, B.D., Joseph, J.M., Roy, N., How, J.P.: Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns. *Autonomous Robots* **35**(1) (2013) 51–76
2. Bandyopadhyay, T., Jie, C.Z., Hsu, D., Ang Jr, M.H., Rus, D., Frazzoli, E.: Intention-aware pedestrian avoidance
3. Basseville, M., Nikiforov, I.V.: Detection of abrupt changes: theory and applications. *Journal of the Royal Statistical Society-Series A Statistics in Society* **158**(1) (1995) 185
4. Bennewitz, M., Burgard, W., Cielniak, G., Thrun, S.: Learning motion patterns of people for compliant robot motion. *The International Journal of Robotics Research* **24**(1) (2005) 31–48
5. Deisenroth, M.P., Huber, M.F., Hanebeck, U.D.: Analytic Moment-based Gaussian Process Filtering. In Bouttou, L., Littman, M., eds.: *International Conference on Machine Learning (ICML)*, Montreal, Canada, Omnipress (June 2009) 225–232



**Fig. 7.** Snapshots of a representative rover/robot interaction, in which the autonomous rover (brown) safely navigates to four goal regions (green ring) while avoiding a small dynamic robot (magenta). The rover's current plan (orange) and history (light orange) are shown, along with the predicted dynamic robot behaviors (blue; darker shades indicate higher likelihoods). Nominal dynamic robot behaviors are marked in gray; behaviors are numbered in first snapshot.



**Fig. 8.** Snapshot of a representative experiment consisting of the autonomous rover navigating to avoid two dynamic robots.

6. Ellis, D., Sommerlade, E., Reid, I.: Modelling pedestrian trajectory patterns with gaussian processes. In: Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on, IEEE (2009) 1229–1234
7. Girard, A., Rasmussen, C.E., Quintero-Candela, J., Murray-smith, R.: Gaussian process priors with uncertain inputs - application to multiple-step ahead time series forecasting. In: Advances in Neural Information Processing Systems, MIT Press (2003) 529–536
8. Grande, R., Walsh, T., Chowdhary, G., Ferguson, S., How, J.: Online regression for non-stationary data using gaussian processes and reusable models. In: Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence (UAI). (2014 (submitted))
9. Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. *Physical review E* **51**(5) (1995) 4282
10. How, J.P., Bethke, B., Frank, A., Dale, D., Vian, J.: Real-time indoor autonomous vehicle test environment. *IEEE Control Systems Magazine* **28**(2) (April 2008) 51–64
11. Ikeda, T., Chigodo, Y., Rea, D., Zanlungo, F., Shiomi, M., Kanda, T.: Modeling and prediction of pedestrian behavior based on the sub-goal concept. In: Robotics: Science and Systems. (2012)
12. iRobot Corporation: iRobot Create programmable robot (2010)
13. Joseph, J., Doshi-Velez, F., Huang, A.S., Roy, N.: A Bayesian nonparametric approach to modeling motion patterns. *Autonomous Robots* **31**(4) (2011) 383–400
14. Kelley, R., Nicolescu, M., Tavakkoli, A., King, C., Bebis, G.: Understanding human intentions via hidden markov models in autonomous mobile robots. In: Human-Robot Interaction (HRI), 2008 3rd ACM/IEEE International Conference on, IEEE (2008) 367–374
15. Luders, B., Kothari, M., How, J.P.: Chance constrained RRT for probabilistic robustness to environmental uncertainty. In: AIAA Guidance, Navigation, and Control Conference (GNC), Toronto, Canada (August 2010) (AIAA-2010-8160).
16. Michini, B., Cutler, M., How, J.P.: Scalable reward learning from demonstration. In: IEEE International Conference on Robotics and Automation (ICRA), IEEE (2013)
17. Park, S., Deyst, J., How, J.P.: Performance and lyapunov stability of a nonlinear path-following guidance method. *AIAA Journal on Guidance, Control, and Dynamics* **30**(6) (November-December 2007) 1718–1728
18. Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A.: Ros: an open-source robot operating system. In: ICRA Workshop on Open Source Software. Volume 3. (2009)
19. Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. The MIT Press (December 2005)
20. Rusu, R.B., Cousins, S.: 3d is here: Point cloud library (pcl). In: Robotics and Automation (ICRA), 2011 IEEE International Conference on, IEEE (2011) 1–4
21. Trautman, P., Ma, J., Murray, R.M., Krause, A.: Robot navigation in dense human crowds: the case for cooperation
22. Trautman, P., Krause, A.: Unfreezing the robot: Navigation in dense, interacting crowds. In: Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on, IEEE (2010) 797–803
23. Vasquez, D., Fraichard, T., Laugier, C.: Incremental learning of statistical motion patterns with growing hidden markov models. *Intelligent Transportation Systems, IEEE Transactions on* **10**(3) (2009) 403–416

24. Zhu, Q.: Hidden markov model for dynamic obstacle avoidance of mobile robot navigation. *Robotics and Automation, IEEE Transactions on* **7**(3) (1991) 390–397